NATIONAL PHYSICAL LABORATORY

SYMPOSIUM No. 10

Mechanisation of Thought Processes

VOLUME II

NATIONAL PHYSICAL LABORATORY

SYMPOSIUM No. 10

Mechanisation of Thought Processes

Proceedings of a Symposium held at the National Physical Laboratory on 24th, 25th, 26th and 27th November 1958

VOLUME II

LONDON: HER MAJESTY'S STATIONERY OFFICE

1959

SESSION 4A

IMPLICATIONS FOR BIOLOGY

Chairman: PROF. J. Z. YOUNG, University College, London

SESSION 4A

PAPER 1

SENSORY MECHANISMS, THE REDUCTION OF REDUNDANCY, AND INTELLIGENCE

by

DR. H. B. BARLOW

BIOGRAPHICAL NOTE

HORACE BARLOW, Physiological Laboratory and King's College, Cambridge. Age 36. Studied Physiology at Cambridge and Medicine at Harvard and London (Univ. College Hosp.). Has worked on various aspects of vision, including eye-movements; spatial properties of receptive fields in the frog's retina; changes in temporal and spatial summation with level of adaptation; and thresholds as signal/noise discriminations. Worked for a year with S. W. Kuffler at Johns Hopkins on changes in retinal organisation in the cat's retina during dark adaptation. Interested mainly in the nervous organization of the visual pathways.

SENSORY MECHANISMS, THE REDUCTION OF REDUNDANCY, AND INTELLIGENCE

by

DR. H. B. BARLOW

SUMMARY

PSYCHO-PHYSICAL and physiological investigations have shown that the eye and the ear are remarkably efficient instruments: consequently the amount of information being fed into the central nervous system must be enormous. After a delay, which may vary from about 100 msec. to about 100 years, this information plays a part in determining the actions of an individual: therefore some of the incoming information is stored for long periods.

The argument is put forward that the storage and utilization of this enormous sensory inflow would be made easier if the redundancy of the incoming messages was reduced. Some physiological mechanisms which would start to do this are already known, but these appear to have arisen by evolutionary adaptation of the organism to types of redundancy which are always present in the environment of the species. Much of the sensory input is not shared by all individuals of a species (eg. stimuli provided by parents, language, and geographical locality) so a device for "learning" to reduce redundancy is required. Psychological experiments give indications of such mechanisms operating at low levels in sensory pathways, and "intelligence" may involve the capacity to do the same at high levels.

In order to exemplify the operations contemplated, a device which reduces the correlated activity of a pair of binary channels is described.

THE usual mechanistic approach to the higher nervous system begins with a consideration of the factors which can be shown to have an immediate effect on the output of the nervous system. The commonest starting point is the simple monosynaptic reflex in which a single sensory input controls a single motor output, as shown diagrammatically in *fig. 1(a)*. The next stage is to elaborate this by taking into account other sensory modalities, inhibition, internuncial neurones, and controlling neurones from elsewhere

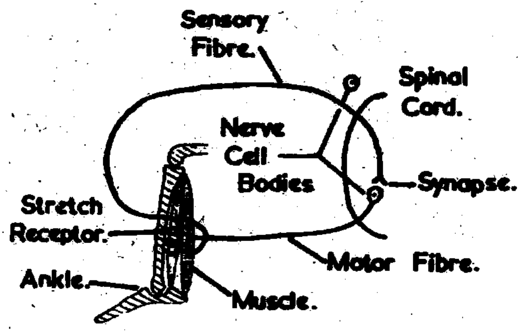

Fig. 1(a)

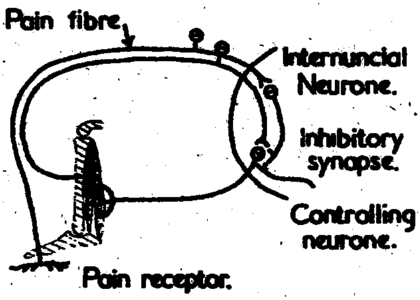

Fig. 1(b)

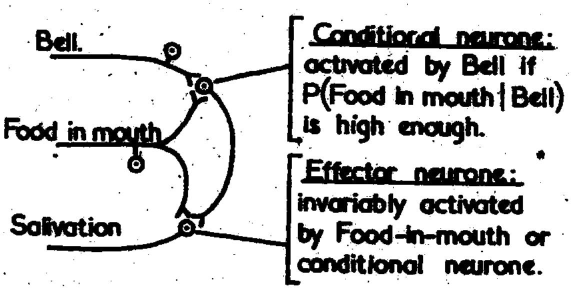

Fig. 1(c)

Fig. 1. Diagram showing approach to higher nervous function from motor (effector) side. (a) monosynaptic stretch reflex; (b) same with addition of internuncial neurones, controlling neurones from other parts of the central nervous system, and inhibition by pain endings; (c) conditioned reflex.

in the nervous system, as shown in *fig. 1(b)*. With all its trimmings this gets one to a stage of complexity perhaps comparable to that of an automatic tracking radar set, or the automatic pilot of an aeroplane. It will show none of the plasticity or adaptability to new surroundings which is characteristic of the higher nervous system, so the Pavlovian conditioned reflex is next introduced. The principle here is that if there are two sensory stimuli (Bell and food in mouth), one of which (food in mouth) always produces a response (salivation), then if they occur jointly with sufficient frequency, the one which, to begin with, did not cause a response, begins to do so (Bell alone causes salivation). This is shown diagrammatically in *fig. 1(c)*, and is perhaps the simplest type of learning behaviour that has been studied in animals, though it has not been investigated in a simple isolated preparation as the diagram might suggest. Uttley (1954, *refs. 22 and 23*) has clarified the principles of operation of such mechanisms and built conditional probability devices which show the same properties of learning and inference.

Now the simple feedback diagram in *fig. 1(a)* has a single input channel, *fig. 1(b)* and *(c)* have two inputs, and Uttley's machine has up to five inputs; but a human brain has something like 3×10^6 sensory nerve fibres leading into it. If it could be supposed that a million or so devices like that of *fig. 1(c)* would deal with the sensory inflow one would be well satisfied with the understanding gained from this approach: but this is not so. The essential operation in a conditional probability device is to measure the frequency of occurrence of *combinations* of activity in the input. Now if the number of binary inputs is increased from two to a million the number of possible combinations is increased from 2^2 to $2^{(\text{million})}$; an arrangement like that of *fig. 1(c)* takes one less far than at first sight appears. I think it follows from this consideration that conditional probability machines cannot be fed with raw sensory information, and the problem of digesting or processing the sensory information entering the brain is an important one. Furthermore, modern electrophysiological techniques are making it possible to record from nerve cells at various levels in the sensory pathways, so this is a problem which is becoming accessible to experimental investigation.

In this paper I have first tried to make rough estimates of the rate at which information flows into the human brain. It is then suggested that an essential step in organising this vast inflow is to derive signals of high relative entropy from the highly redundant sensory messages. For this something similar to the optimal codes discussed by Shannon (1949, *ref. 19*) needs to be devised for the sensory input, and the steps required to do this are considered. Finally, a modified form of such recoding is proposed, some evidence that it occurs is brought forward, and it is suggested that the idea may be extended to cover some of the processes going on in consciousness and called reasoning or intelligence.

1. THE SENSORY INFLOW

(a) Properties of Nerve Fibres

We are equipped with sensory instruments of astonishing sensitivity and versatility which supply information about the environment to the central nervous system. This information is carried along nerve fibres, and since a good deal is known about what these fibres can and cannot do, one can derive an approximate upper limit to the rate at which information enters the brain. If the simple assumptions are made that (i) the maximum frequency of impulses is 700/sec, and (ii) in 1/700th. sec a nerve can only be used to indicate the presence or absence of an impulse, then the maximum rate at which it can transmit information is 700 bits/sec. Mackay and McCulloch (1952, *ref. 16)* point out that the nerve might be used more efficiently if, instead of detecting the presence or absence of an impulse, the intervals between impulses are used to convey information. Using such pulse interval modulation, and assuming (i) accuracy of estimation of intervals of 0.05 msec, (ii) a minimum interval of 1 msec, they give the maximum capacity as 2880 bits/sec. This would require a mean frequency of 670 impulses/sec, but at a mean frequency of 50/sec, such pulse interval modulation still allows 500 bits/sec to be transmitted. These figures are actually too low, because Mackay and McCulloch incorrectly assumed that the optimum distribution of intervals was uniform instead of exponential: however, if the other assumptions are granted, they show clearly that a single nerve fibre could be used to transmit information at a rate well above 1000 bits/sec.

The total capacity of the sensory inflow appears to be above 3×10^9 bits/sec, but it is certain that nothing like the full capacity is utilised. The mean frequency of impulses must be far below the optimum; peripheral nerves appear to use pulse frequency rather than pulse interval modulation, so that there will be high serial correlations between the values of intervals; furthermore, there are generally considerable overlaps in the pick-up areas of neighbouring fibres, which are therefore bound to show correlated activity. Finally, the figure for the performance of a nerve fibre given above might be approximately true for the large diameter fibres, but those of smaller diameter, which make up a large fraction of the total number, must have a smaller capacity. It would be pure guesswork to try to allow for these factors, but one can get indications of the utilised capacity from two other sources.

(b) Sensory Ability

Jacobson (1950, 1951 *refs. 13,14*) has made estimates of the informational capacity of the ear and the eye. For the ear he calculated 50,000 bits/sec from the number of discriminable pitches (about 1450), the number of discriminable intensities at each pitch (average about 230), and the time required to make such discriminations (1/4 sec). This does not make any allowance for masking - the observed fact that the presence of one tone interferes with the perception of other tones. Jacobson calculated that this would reduce the information capacity by a factor of about six, bringing it down to 8,000 bits/sec. Now there are 30,000 nerve fibres from the ear, so each fibre must carry an average of about 0.3 bits per sec.

For the eye he calculated from published data of central and peripheral acuity that there were 240,000 resolvable elements in the visual field (he seems to omit a factor of two in the integration, but this is perhaps compensated by the rather high figure for acuity which he uses). He supposes that each element can be discriminated at two intensities, with an average temporal resolution of 1/18 sec. These figures give 4.3×10^6 bits/sec. In the optic nerve there are just under a million fibres, so about 5 bits/sec are conveyed on the average by each fibre.

These are crude estimates. For instance, no account has been taken of colour discrimination, or of the ability to localise a sound by binaural effect and judge depth by stereoscopic vision. Nevertheless, they are probably of the right order of magnitude and they are probably good enough to justify the claim that optic nerve fibres carry much more information than those of the auditory nerve. This may be significant and will be referred to later.

These figures suggest that total sensory inflow along the three million sensory fibres is rather under 10^7 bits/sec.

(c) Communication bandwidths

The capacity of the communication channels engineers need to transmit auditory and visual signals is clearly related to the capacity of the sensory pathways. Engineers, in the interests of economy, may be expected to try to use the narrowest bandwidths which will satisfactorily load up the sense organs involved, and recipients may be expected to insist that such satisfactory loading is not too far short of normal loading.

Ten k.c. bandwidth at 40 d.b. signal noise ratio give a good quality auditory signal, and has a capacity of 133,000 bits/sec. This is more than ten times Jacobson's final figure for the capacity of the ear (8,000 bits/sec), and the discrepancy is presumably due to (i) the transmission of relative phases of the frequency components, which gives information not utilised by the ear - at least in the type of discrimination taken account of by Jacobson; (ii) the failure of the engineer to exploit the loss of efficiency of the ear which results from masking.

A satisfactory 400 line television picture requires three megacycle bandwidth at about 10 d.b. signal-noise ratio, and this corresponds to 1.2×10^7 bits/sec. One is much more aware that such a television picture falls short of one's normal visual signals than one is in the case of a

10 k.c. 40 d.b. auditory signal because it does not fill the visual field, and lacks detail and colour, but it is still more than double Jacobson's estimate of the eye's capacity. In this case the most notable matching errors are the failure to exploit (i) low peripheral acuity of the eye, (ii) reduced temporal and spatial resolving power in low intensity regions of the image.

Engineers seem to require 5 - 10 bits/sec channel capacity per nerve fibre to load up our sensory pathways, but the discrepancies between this figure and those obtained from direct estimates of sensory abilities can probably be attributed to poor matching.

(d) Time of storage

Not only is the input to the nervous system enormous, but some, at least, of the messages received are stored for very long periods. Most people would agree that sensory impressions can be recalled after a lapse of, say, 70 years, and sometimes a person can produce objective evidence of the accuracy of his recollections. In addition there are, of course, many sensory impressions which cannot be recalled, but which have, none the less, left their mark: we do not remember the successes and failures by which we acquired the correct usage of 'yes' and 'no', but this correct usage is often retained beyond the retiring age. If one allows for fifty years of waking life, the total sensory input is something like 10^{16} bits. Complete storage of all this information is neither likely to be possible nor, of course, is it what is needed.

(e) Fate of Sensory Information

The rest of this paper is about a suggested plan of storing and displaying this enormous sensory input, but one must first have some idea of the use that is made of the sensory information and the neural equipment which is available for dealing with it. According to Craik (1942, *ref. 4*) the sensory information is used to build up a model of the external world which provides a basis for determining what course of action is most likely to lead to the survival of the individual and his species. That is a brief answer to the first question, and it also gives the answer to another fact which might otherwise be puzzling. A man can only make decisions on the basis of sensory information at a maximum rate of about 5 to 25 bits/sec. (Hick, 1952, *ref. 11* Quastel, 1956, *ref. 18*): why, then, does he need a sensory input of 10^7 bits/sec.? Craik's answer would probably have been that the greater the sensory input the more complete and accurate the model, and hence the surer its basis for planning survival.

The question of the equipment available can also, because of our ignorance, be answered briefly. There are some 10^{10} interconnecting nerve cells in the central nervous system, and quite a large proportion of them must be available for the task of dealing with the sensory input and building up the model. We are only beginning to determine the properties of these cells; it has been known that their long processes transmit information as all-or-none impulses for more than fifty years, but how information is stored is not yet understood. In what follows I shall be talking about *what* the nervous system does rather than *how* it does it, so our ignorance of the method of storage of information is not too serious. The problem might be discussed abstractly, but for the sake of a definite model one can think of each nerve cell having "excitation laws" which determine the conditions under which it becomes active, and suppose that these laws can be changed so that it becomes active in response to a different set of patterns of activity in the nerve cells in contact with it. The excitation laws for all the neurones would then form a store of information and the current display would consist of the pattern of nerve cells which are actually transmitting impulses down their long processes at any given moment.

With this model in mind the problem is: what should the excitation laws of the neurones be, and how should they be alterable, in order that the display of activity shall help the individual and species to survive in the situation giving rise to the current sensory input? To avoid basing the argument on uncertain preconceptions of what the brain does, one could put it in more general terms in this way. The barrage of nervous impluses reaching the nervous system seems to be unmanageably large; how should a selection of this activity be made for current display and future reference?

2. ORGANISATION OF THE SENSORY INPUT

The proposition is that the initial selection is performed according to those statistical properties of the past sensory messages which determine how much information particular impulses convey. It is supposed that the sensory messages are submitted to a succession of re-coding operations which result in reduction of redundancy and increase of relative entropy of the messages which get through. Ideally one might imagine that an optimal code is constructed, so that the output, or "display" of current input, has no redundancy, relative entropy 1, and carries all the information of the input. This ideal obviously cannot be reached, but the re-coding operations are supposed to tend towards the ideal: that is, outputs are derived from the input, which have high relative entropy and carry as much of its information as possible.

Shannon has shown that it is possible in principle to obtain near optimal coding if a sufficient number of messages of a given length have occurred to give knowledge of the statistical structure of the messages, and if delays are permitted between input and output. Fano and Huffman (1953, *ref. 12*) have described procedures for constructing such codes. The first steps are to define what shall constitute a single message and then to measure the frequency of occurrence of all possible messages of this class. Clearly the class cannot be the whole of the sensory input to the brain up to a particular moment, for this message has only occurred once. The input must be sub-divided in time, and first consider the operation required to re-code messages of duration, say, one second. The capacity of the input channel has been shown to be about 3×10^9 bits/sec. which corresponds to $10^{(\text{thousand million})}$ possible messages per second. If one takes account of the restrictions which reduce the utilised capacity to some 10^7 bits/sec., and considers messages of one-tenth second duration, there are still some $10^{300,000}$ possible messages. It would clearly be hopeless to devote neural equipment to the counting of each possible message, for it is highly improbable that any single message will be exactly repeated and most of such equipment would be unused at death. This is, essentially, the same difficulty that was levelled against the idea that conditional probability devices could be served with unprocessed sensory data, but when one considers optimal coding there is a possible solution. Because the code is reversible, no information is lost by re-coding small sections of the sensory input independently, and such preliminary re-coding will enable the whole message to be passed down a channel of smaller capacity, and thus facilitate subsequent steps.

The idea is best illustrated by considering the order in which different types of redundancy might be encountered, and eliminated, during the successive re-coding operations. First there is the very large amount which results from the inefficient utilisation of peripheral nerve fibres. Looking only at the nerve impulses as they arrive, it would be found that impulses occurred at different mean rates in different fibres and in all of them at rates well below the optimal frequency for information transmission. This type of inefficient utilisation of a set of communication channels is a form of redundancy, but for reasons discussed later (Section 4) it may be less important to eliminate than other forms: for the moment one can consider the capacity of a nerve fibre as determined, not by maximum frequency of impulses, but by the mean frequency at which they occur.

Next, still looking only at the impulses as they reach the central nervous system, it would be found that impulses do not occur completely at random in time but tend to follow one another in sequences and bursts: the first re-coding operation might be a mechanism which reduced the serial correlations so that the same amount of information was carried by fewer impulses. In addition it would be found that certain groups of nerve fibres tended to become active at the same time. These would be fibres whose receptive fields on the sensory surface overlapped, so that this particular form of redundancy results from the anatomical properties of fibres and sense organs, just as the serial correlations in time result from the fact the intensity of a stimulus is coded as frequency of impulses at the sense organs.

These first steps, then, would reduce the orderliness in the sensory messages which results from characteristics of the sensory apparatus. But if this orderliness can be eliminated, so can that resulting from the characteristics of the environment which is providing these stimuli. For instance, it will often happen that a stimulus covers more than a single point on the sensory surface and therefore causes activity in a group of fibres larger than those whose receptive fields overlap. Advantage could be taken of this to reduce the number of impulses required to convey information about such a stimulus. Again, a stimulus will often be moved across a sensory surface causing excitation in sequences of nerve fibres. Such repeated, ordered, sequences of activity would be a form of redundancy which could be reduced by suitable re-coding. In fact, any pattern of stimuli which represents a departure from complete randomness - such as simultaneous stimuli at different points on the sensory surface, stimuli which are maintained for long duration of time, ordered sequences or cycles of stimuli - present an opportunity of reducing the magnitude of the sensory inflow by suitable re-coding. It is clear that many of the complex features of our environment will come into this category. For instance, the stimuli which result from an animal's parents or its habitat are repeated frequently, and economies could be effected by reducing the space in the sensory representation occupied by these familiar stimuli and allowing more space for the infrequent and unexpected stimuli.

It is suggested, then, that the processing or organisation of sensory messages is carried out by devising a succession of optimal or near-optimal codes adapted to the messages which have been received. In the early stages the total inflow will be sub-divided into many small sections, presumably taking in each section the messages coming along neighbouring fibres during a short interval of time. In the later stages the coded outputs will be re-mixed, possibly with the addition of delayed inputs (as utilised by Uttley in conditional probability devices) to allow detection of movement and other ordered sequences of activity, and then will be sub-divided again into small sections. Thus in the later stages the nerve messages being re-coded may be derived from more and more remote parts of the sensory inflow and may also come from sensory stimuli more and more separated from each other in time of occurrence. It will be seen that at each stage storage of some of the sensory information is required in order to construct the optimal code, and thus the code itself forms a kind of memory.

Now the idea that our brains detect order in the environment is not new. Empiricist philosophers have talked of percepts being associated sense impressions, and of causality corresponding to invariant succession of sense impressions. Behaviourists have emphasised the importance of association, and Gestalt psychologists talk of ordering sensation according to certain schemata (though here there seems to be some confusion as to whether the ordered schemata are derived from sensations or imposed upon them). Thus the fact that our higher centres are much concerned with the redundancy of the sensory messages has often been pointed out, but two aspects of this fact have not, I think, been so widelv recognized. First, the detection of redundancy enables the sensory messages to be represented or displayed in a more compact form; and second, the reduction of redundancy is a task which can be subdivided and performed in stages. *Figure 2* shows diagrammatically how the suggested scheme of storage and display compares with more othodox representations of memory and consciousness. It will be seen that in the present scheme a large part of the storage of information occurs before the display - that is before the level of re-coding which might correspond to conscious awareness of sensory stimuli. The re-coding is supposed to continue at conscious levels, so some of the information reaching consciousness is also stored, but this would only be sufficient, first, to enable the process of building up the code to continue, and second to enable "useful" association to be made between motor acts and features of the current sensory input (e.g. between salivation and bell).

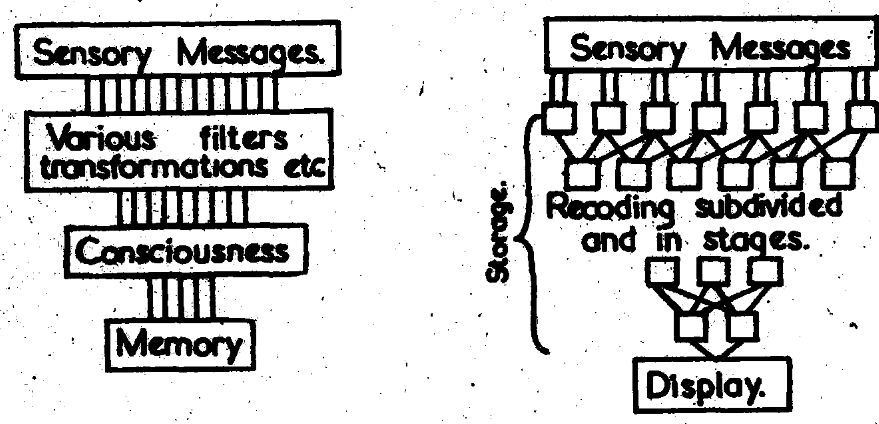

Fig. 2. Diagram contrasting memory *after* consciousness in orthodox scheme with storage *before* display in optimal coding scheme.

It seems a help to consider the processing of sensory information as optimal or near-optimal coding for two reasons. Practically, it enables the subject to be approached along the firm path of sensory physiology instead of through the shifting sands of conscious introspection and philosophy. And conceptually it shows a way in which complete mental acts, which seem appalling in their complication and perfection, may be sub-divided into a succession of much simpler operations; this is clearly a prerequisite for gaining an understanding of the physiological basis of mental function.

It is worth noting that the possibility of sub-division rests on Shannon's proof of the possibility of near-optimal coding; if the early transformations of the sensory information were not reversible, redundant features which are detected later might be lost: and if the earlier transformations did not increase the relative entropy of the messages, they would not facilitate the detection of higher order redundancy.

3. DESIRABILITY OF OPTIMAL CODING

In the last section an outline scheme for dealing with the enormous sensory inflow was suggested. In this section some reasons for the desirability of optimal coding are put forward. It will be argued that it is desirable on the grounds of accessibility, stability, and economy, and because it requires storage of information sufficient to form a model of the animal's environment. Of course, such arguments for its desirability are not sufficient reasons for believing that it actually occurs.

(a) Accessibility.

Optimal coding will improve the accessibility of information in two ways. First, the capacity of the display required for the current sensory input will be decreased. This simplifies the task of finding useful associations just as reducing the size of a haystack simplifies the task of finding needles. The second way is less obvious. In messages of high relative entropy, the probability of a given message occurring is close to the product of the probabilities of the individual signs which make it up. Now a dog feeds once or twice a day, and when looking for sensory correlates of salivation it would not be worthwhile to search among combinations of individual signs whose probability of joint occurrence was so low that they would be expected only, say, once a week, nor amongst those whose probability was so high that they would be expected, say, once an hour. If the input to a conditional probability device is known to be of high relativ entropy, great economies of design are possible.